

Evolutionary or Revolutionary? Applied Mathematics for Exascale Computing

SIAM Annual Meeting, Chicago, IL
11 July 2014

Jeff Hittinger



LLNL-PRES-656899


This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC



**If you had asked me several years ago
about Exascale Computing...**


meh.

We lack the computing power to tackle Grand Challenge Science problems




Combustion

- High-pressure, turbulent reacting flow
- Complex moving geometry
- Multiphase: fuel injection and soot
- Stochasticity
- Optimal engine design



Climate

- Coupling atmosphere, oceans, ice sheets, land mass, biosphere
- Global to microscopic
- Catastrophic rare events
- Extreme weather patterns
- Assessments for policy



Materials

- Transient mesoscale behavior of new materials
- Search for novel, optimal materials
- Model from nanometers to microns, femtoseconds to minutes

Need (at least) exascale computing resources

What is an exascale-class machine?

	ASCI Red	Road Runner	K Computer	Sequoia		Exascale
Year	2000	2008	2011	2012		2023
Peak (Flops)	1.3e12	1.7e15	11.3e15	20.1e15		1.2e18
Linpack (Flops)	1.0e12	1.0e15	10.5e15	16.3e15		1.0e18
Total Cores	9,298	130,464	705,024	1,572,864		1e9
Processors	9,298	12,960(6,912)	88,128	98,304		1e6
Cores/Proc	1	9(2)	8	16		1e3
Power (MW)	0.85	2.35	9.89	7.9		~20

Adapted from B. Harrod, "DOE Exascale Computing Initiative Update," Aug 15, 2012

Power has become the dominant constraint

Based on current technology, scaling today's systems to an exaflop level would consume more than a **gigawatt** of power, roughly the output of Hoover Dam

– 2012 ASCAC Report “The Opportunities and Challenges of Exascale Computing”




Using commodity hardware:
Exascale machine: \$100B
Annual Power Bill: \$1B
Phenomenal science: Freeless

[Hoover Dam at Night](#), Tex Roy Bean, [CC BY-SA 3.0](#)


Power is also driving architecture changes

- *Power densities limit clock speeds*
- *More cores and specialized accelerators*
- *Data motion costs on-chip and off-chip*
- *Volatile memory (DRAM) is power-hungry*



Original data collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond and C. Batten.
Dotted line extrapolations by C. Moore. From C. Moore, "Data Processing in Exascale-Class Computer Systems," Salishan, 2014

Exascale computing introduces several fundamental challenges



Extreme Concurrency

- Processing units ↑
- Bulk-synchronous will not scale
- Concurrency ↑
- Synchronization ↓
- Communication ↓
- Dynamic task parallelism



Limited Memory

- Memory gains less than processing
- Memory/core ↓
- Minimize memory usage
- Deeper , heterogeneous memory hierarchies



Data Locality

- Transfer gains less than processing
- Bandwidth/core ↓
- Energy and time penalties for data motion
- Greater need for data locality
- Reduce data transfers



Resilience


- Massive number of components: hard faults ↑
- Running closer to threshold voltage: soft faults ↑
- Bulk-synchronous checkpoint restart is dead

Will Mathematics for Exascale be...



Mick Tsikas, Reuters

OR



DOE ASCR chartered an Exascale Applied Mathematics Working Group

Charge

Identify:

- gaps in thinking about exascale
- new algorithmic approaches
- new scientific questions
- a more holistic approach

Team


Jack Dongarra*
John Bell
Luis Chacon
Rob Falgout
Mike Heroux

Jeff Hittinger*
Paul Hovland
Esmond Ng
Clayton Webster
Stefan Wild

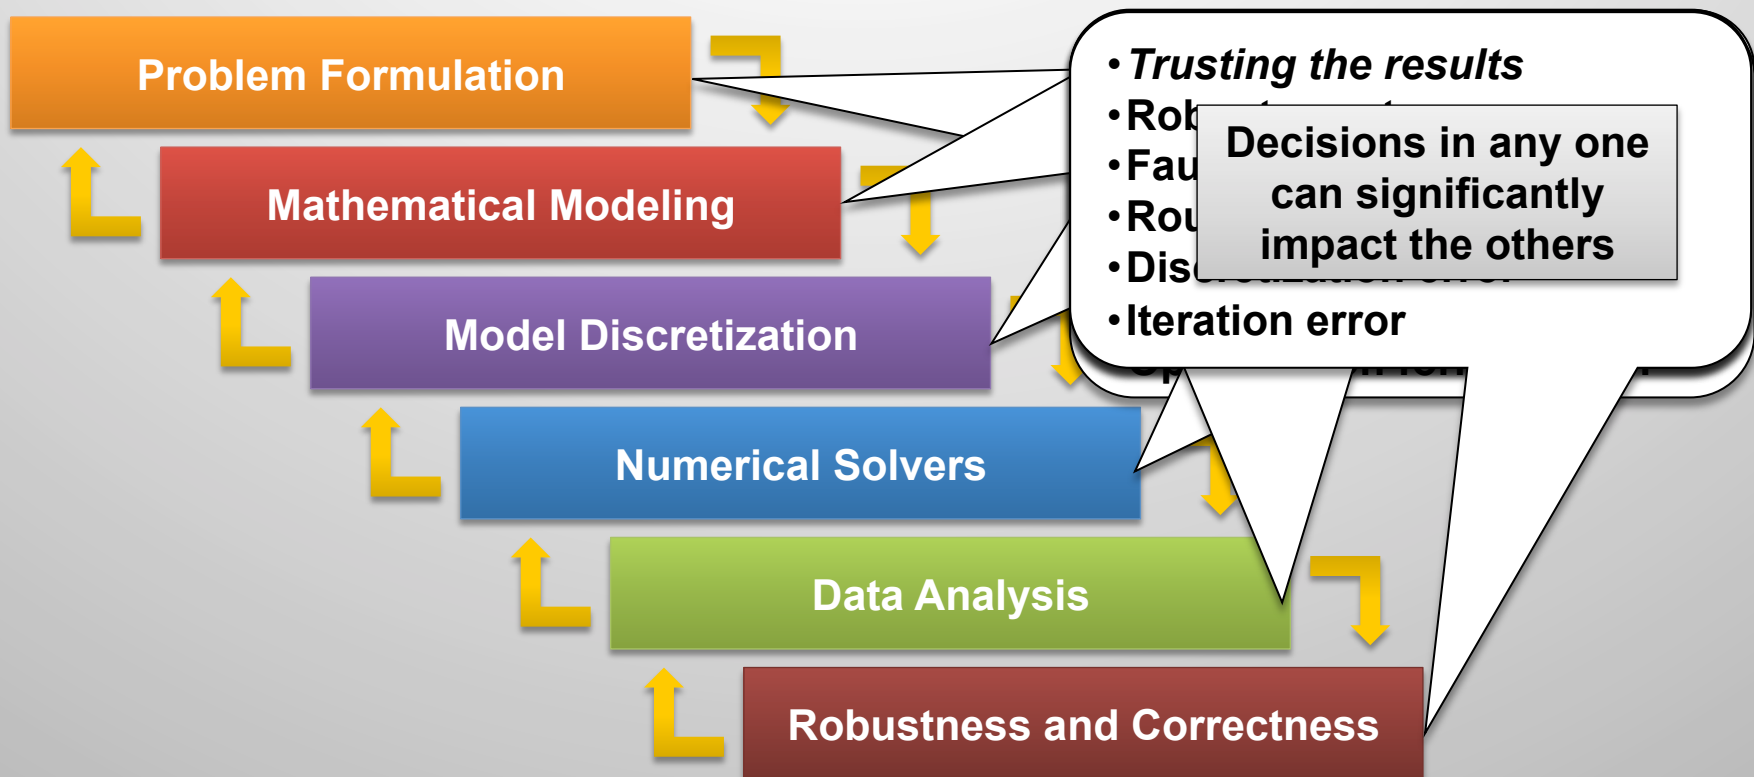
*co-chairs

Process

- **Community Workshop (Aug 2013)**
- **Fact-finding teleconferences**
- **Grand Challenge reports**



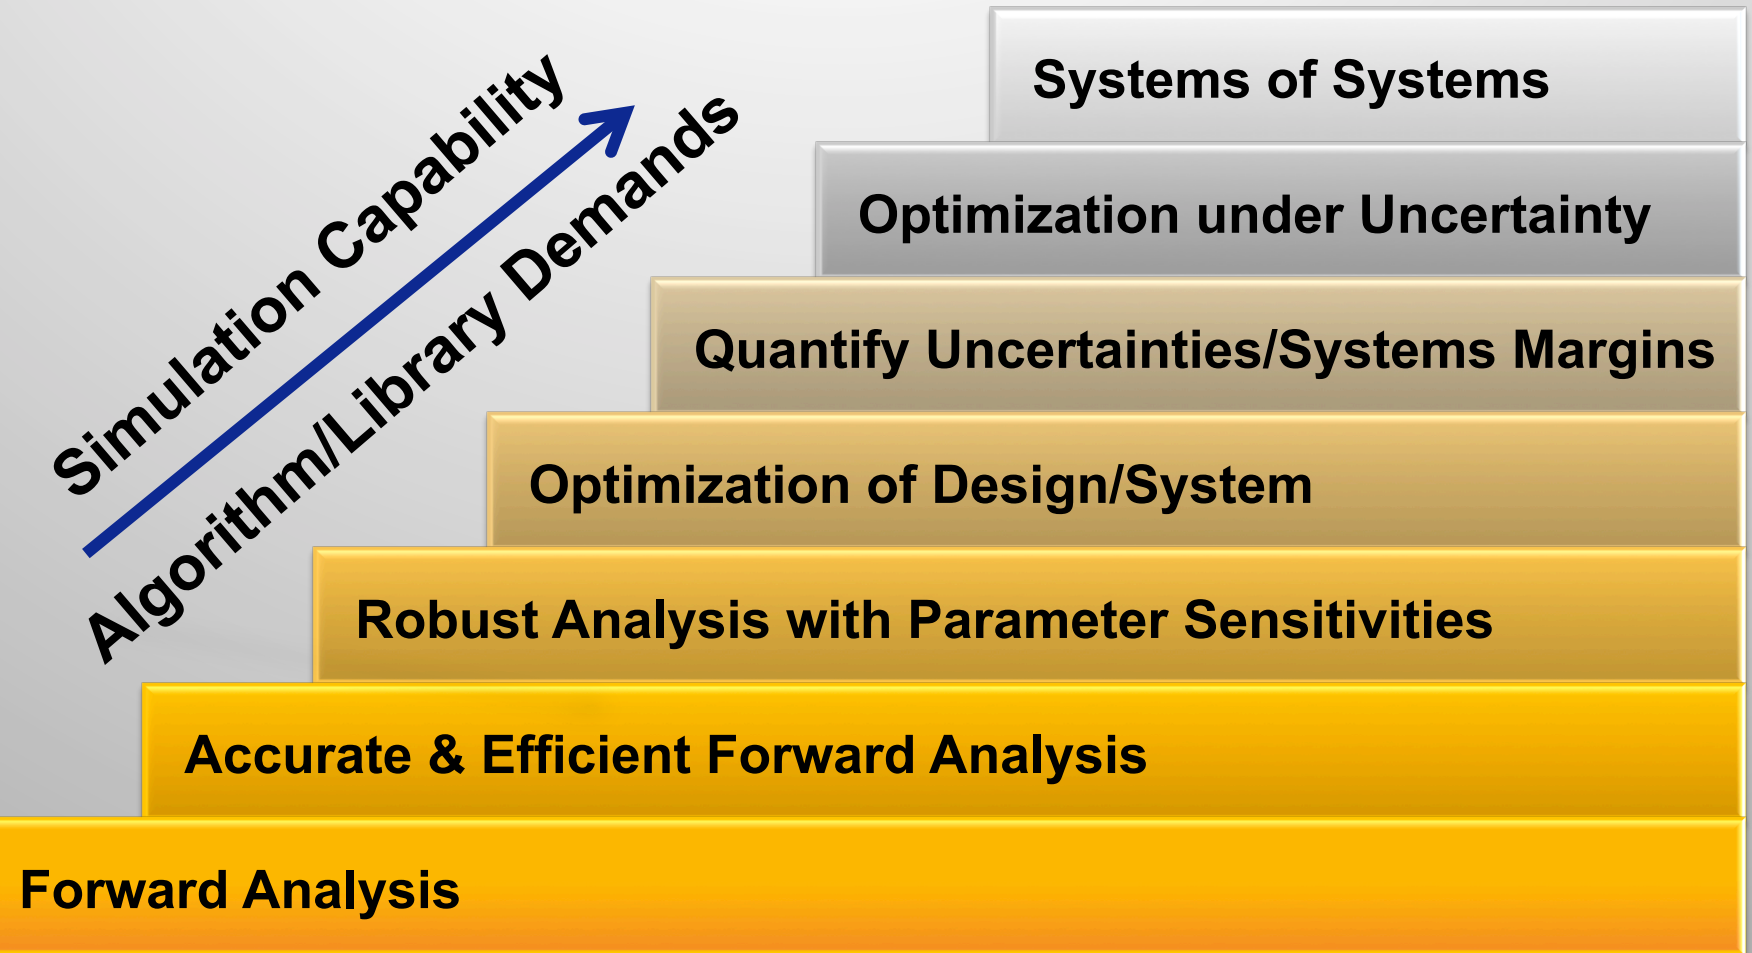
An organizing principle we used was the concept of the *Mathematics Stack*



Areas outside of this conceptual organization:

- Optimization and optimal control for system management
- Discrete mathematics and graph analysis
- Finite state machines and discrete event simulation

Problem Formulation: A dramatic potential to change the questions we ask




Oberkampff, Pilch, Trucano, SAND2007-5948, SNL, 2007

Mathematical Modeling: In forward simulation, we must consider new models

- Can we model additional physics?
- How else can we model the problem?
- Do some models expose more concurrency?
- Scale-bridging models
 - Hierarchical representations
 - Coarse-graining
- Particle vs. continuum

We must respect the physics!




Mathematical Modeling: Uncertainty quantification plays a larger role at exascale

We must be clever in combating the curse of dimensionality

Performance Increase 3D FEM Nonlinear Diffusion

Phipps, Edwards, Hu, Webster, Equinox project, ASCR XUQ

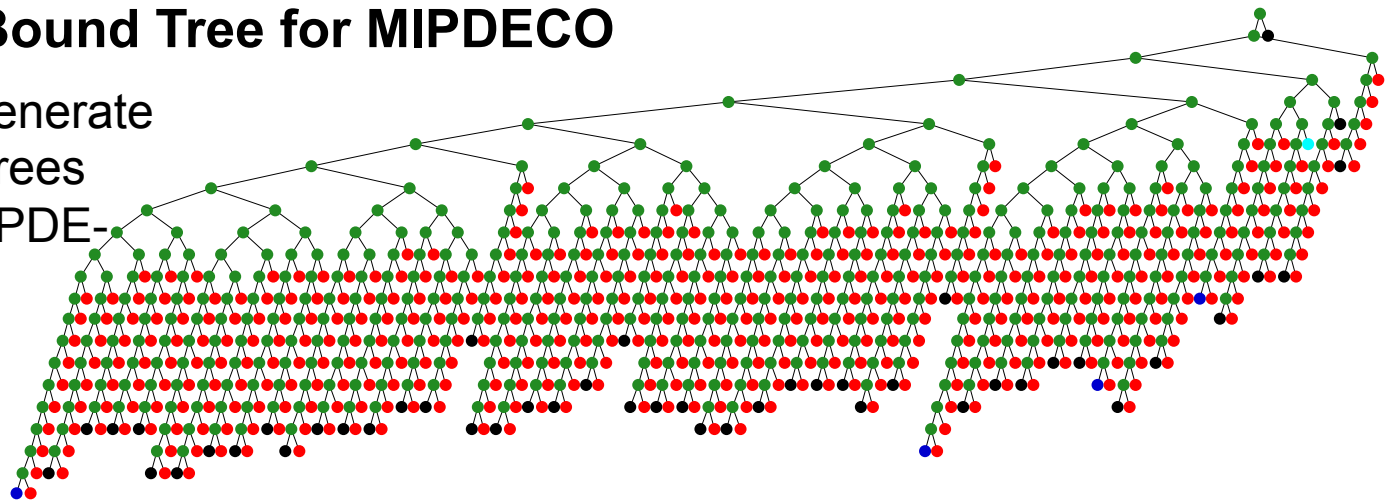


- Adaptive hierarchical methods
- Advanced multilevel methods
 - Model hierarchies
 - Stochastic hierarchies
- Architecture-aware UQ
- Adaptive and robust methods for fusing computation and experimental data

Mathematical Modeling: Exascale will enable the solution of new optimization problems

Branch and Bound Tree for MIPDECO

- MIPDECOs generate huge search trees
- Each node is PDE-constrained optimization




[Leyffer & Mahajan]

- Concurrent-point methods
- Mixed-integer, simulation-based, and global optimization
- Multi-fidelity hierarchies
- Robust optimization and optimization under UQ
- Optimal design and coupling of experiments

Discretization: Partitioned algorithms will play an important role




- Partitioned algorithms in:
 - Models, equations, and operators
 - Spatial (FSI)
 - Temporal (multimethod, multirate)
- Need better coupling strategies
 - High-order
 - Consider splittings based on strength of coupling
 - Compatible interface treatments
 - Nonlinearly converged strategies
- Stability, consistency, and accuracy



Source: J. Banks and W. Henshaw


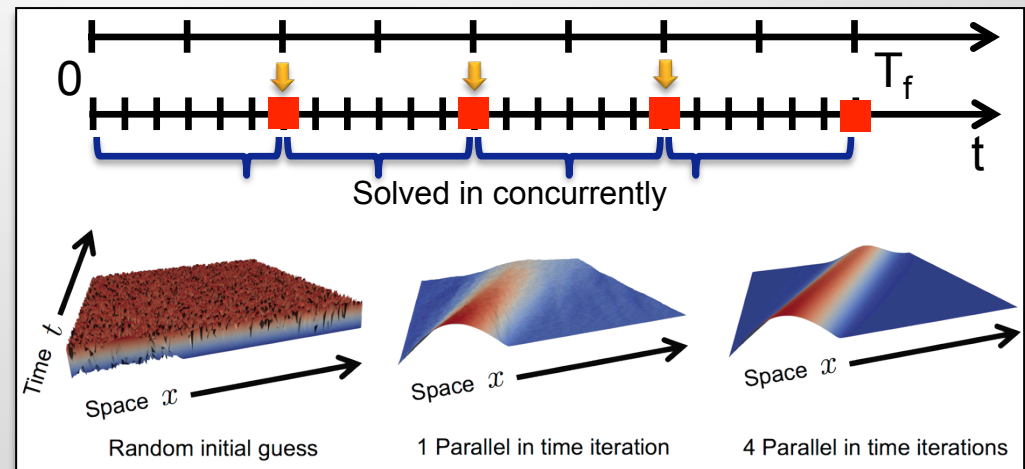
Discretization: It is expected that high-order discretizations will become dominant

- High-order discretizations
 - High arithmetic intensity
 - Maximize on-node performance
 - Robustness?
- Adaptivity
 - Mesh
 - Model
 - Discretization/order
- Scalable computational geometry and mesh generation



Discretization: Overcome sequential bottleneck of time integration


- Parallel-in-time
 - Hierarchy of representations of varying fidelity
 - Iterative time advancement
- Research issues:
 - Optimal convergence
 - Chaotic systems
 - Oscillatory systems
 - Hyperbolic systems



J. Schroder et al., hypre project

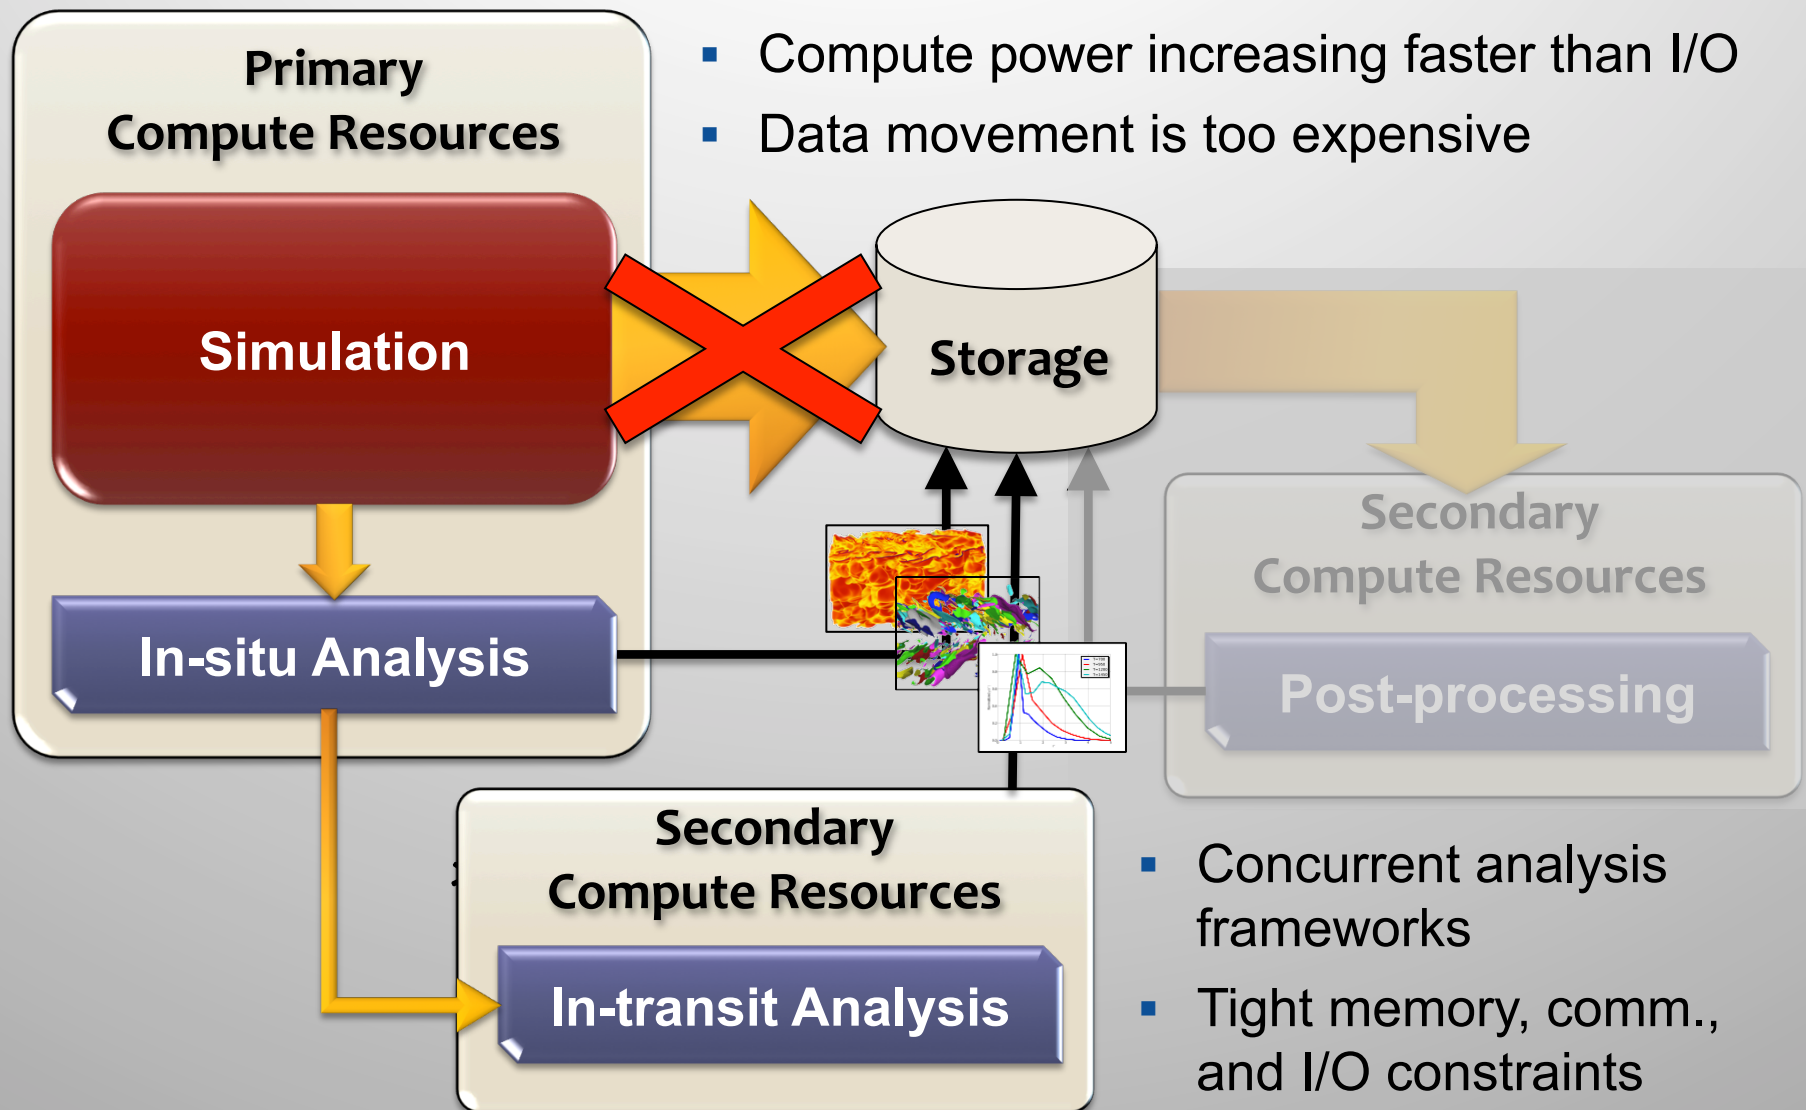
Scalable Solvers: In solving the discrete system, numerous topics must be addressed

- Communication-avoiding
- Synchronization reduction
- Data compression
- Multiple-precision
- Randomization and sampling
- Adaptive response to load imbalance
- Scheduling and memory management
- Autotuning algorithms
- Energy-efficient algorithms



Example: Timing comparison on 100x100x100 7-point Laplacian stencil [E. Chow and A. Patel]

Data Analysis: Understanding the results




Data Analysis: Concurrent analysis assumes a priori knowledge of the features of interest

- Feature-Aware in situ transformations


- Statistical
 - Principal component analysis
 - Isomap
 - Locally linear embeddings
- Segmentation-based
 - Image recognition
 - Merge trees: topology, vorticity, etc.
- Application-specific features

- Memory and compute-efficient

- Sub-Linear algorithms
- Streaming: progressive multi-resolution




E. Balasubramanian et al., Science, 2002



J. Bennett et al., IEEE T. Vis. Comp. Graph., 2011

Resilience and Correctness: Trusting the results in the presence of faults

- Resilient programming models:
 - Skeptical
 - Relaxed bulk synchronous
 - Local failure, local recovery
 - Selective reliability
- Algorithm-Based Fault Tolerance
 - Use properties of models and algorithms to detect (good) or be insensitive (better) to faults
 - Understanding how random faults alter convergence



Data from M. Heroux, M. Hoemmen, K. Teranishi

Resilience and Correctness: Dynamic adaptation impairs determinism



- Reproducibility and verification techniques rely on determinism
- Can we justify cost of enforcing determinism?
- Should we interpret reproducibility and verification statistically?
- Analysis to understand the variability of deterministic algorithms

Mathematics for Exascale System Software



- Autotuning as derivative-free optimization
- Adaptive runtime systems as optimal control
- Mathematically grounded scheduling
- Stochastic performance models

This afternoon, we will discuss four of these areas in more depth

MS 139:

**Opportunities in Applied Mathematics Research for Exascale Computing
Salon 7 – 3rd Floor**

**4:00-4:25 Hierarchical Multilevel Methods for Exascale Uncertainty
Quantification and Optimization**

Clayton G. Webster and Stefan Wild

**4:30-4:55 Mathematical Modeling and Discretization for Exascale
Simulation**

Luis Chacon


5:00-5:25 Discrete Solvers at the Exascale

Esmond G. Ng

5:30-6:00 Resilient Algorithms and Computing Models

Franck Cappello

Evolutionary or revolutionary? A Punctuated Equilibrium perspective for HPC evolution



Transitions may be rapid, but continuity with the past is maintained

Math is the DNA of computing that provides the common thread for (r)evolution

Some approaches may become extinct

Some approaches will adjust and continue

Some disfavored approaches will gain importance

Some dominant approaches will lose importance

Some new approaches will be created

It is unlikely that we will discard the 400+ year legacy of the scientific revolution and begin anew in only a decade

Structure of DNA, CC0 1.0

It's the end of the world as we know it... and I feel fine

It's an opportunity to solve challenging problems



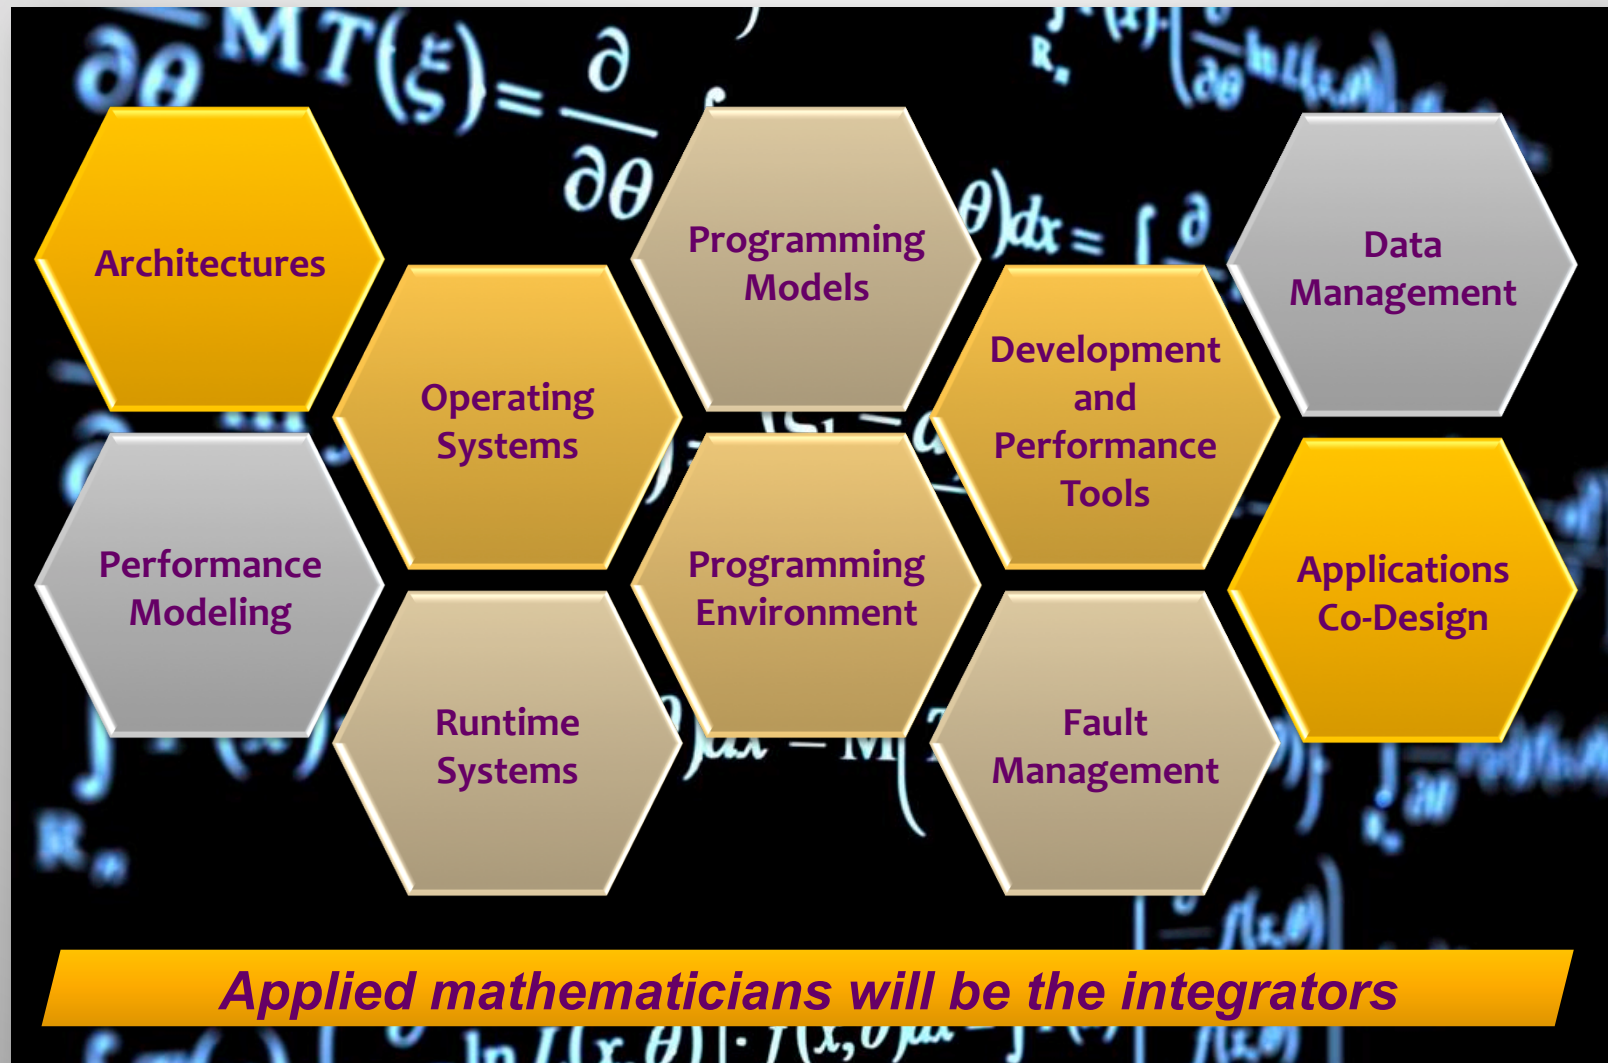
Don Davis, <http://www.donaldedavis.com/BIGPUB/BIGIMPCT.jpg>, CC0



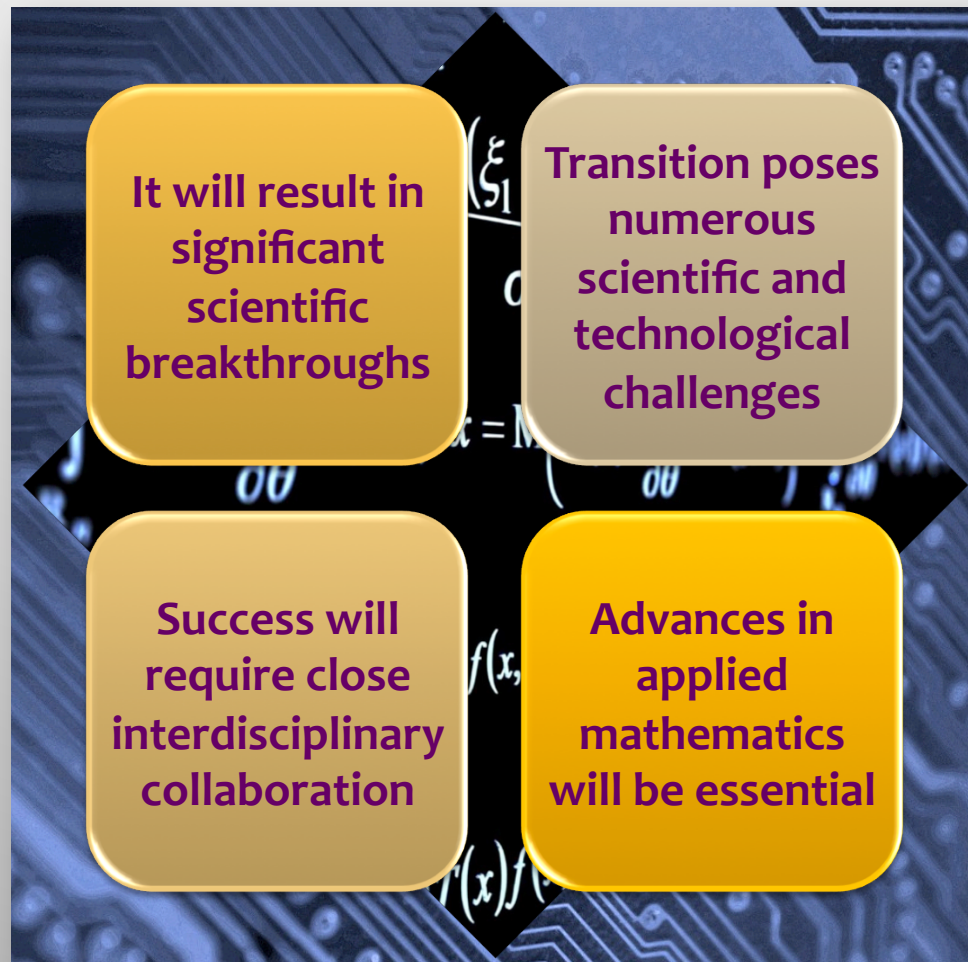
What will emerge?

[Wolpertinger](#), Rainer Zenz, [CC-BY-SA 3.0](#)

The applied mathematics community must work with others to address the challenges of exascale




Exascale computing will allow us to compute in ways that are not feasible today



Many additional resources are available

Exascale Mathematics Report

<http://science.energy.gov/ascr/news-and-resources/program-documents>



Exascale Mathematics Working Group Website

- White Papers
- Workshop presentations
- Background information

<https://collab.mcs.anl.gov/display/examath/Exascale+Mathematics+Home>

DOE Grand Challenge Science Reports



<http://science.energy.gov/ascr/news-and-resources/workshops-and-conferences/grand-challenges>

